THE OFFICE OF ENTERPRISE TECHNOLOGY STRATEGIES

Statewide Technical Architecture

Implementation Guidelines:

Red Hat Enterprise Linux

STATEWIDE TECHNICAL ARCHITECTURE

Implementation Guidelines:
Red Hat Enterprise Linux

Revised Date:		Version:	1.0.0
Revision Approved Date:			
Date of Last Review:	August 28, 2003		
Date Retired:			
Architecture Interdependencies:			
Reviewer Notes: Published on August 29, 2003			

Office of Enterprise Technology Strategies
PO Box 17209
Raleigh, North Carolina 27699-7209
Telephone (919) 981-5510

Introduction

The intent of this document is to outline the implementation guidelines that the State of North Carolina has developed to ensure uniform and consistent implementations of Red Hat's Enterprise Linux (RHEL) operating system (OS) in accordance with the State of North Carolina's Statewide Technical Architecture (STA).

Detailed aspects of RHEL installation procedures will not be covered in this document. There are resources listed in the following sections that outline further support mechanisms for gaining greater depth and knowledge about RHEL. Also, Red Hat has published prescriptive architecture guidance documents that are available on the Internet at http://www.redhat.com/apps/support/ that cover the topics mentioned in this document in much greater detail.

The key goal of this document is to outline implementation guidelines that when followed by the system administrator will lead to a well-designed RHEL solution that has the flexibility to grow with changes in technology and can be maintained in an efficient, effective, and secure manner, which is a fundamental principle of the North Carolina Statewide Technical Architecture. While Red Hat does offer a desktop version of their operating software, this implementation guideline will limit its scope to the server level operating system and associated application offerings of Red Hat.

Red Hat Enterprise Linux was selected for this implementation guideline because of use trends observed throughout state agencies. While other commercially distributed open source operating systems exist, there deployment in state government has not been observed to the level of Red Hat.

This document in no way indicates or implies a preference in a vendor or technology by the State. It is provided to ensure that agencies implementing Red Hat Enterprise Linux have the resources necessary to successfully deploy their systems in the most advantageous manner to both the agency and the State of North Carolina.

Red Hat, Inc.

Marc Ewing and Robert Young founded Red Hat in 1993, not far from the campus of North Carolina State University. Red Hat began as one of hundreds of companies making a business of assembling, testing, and packaging an all-in-one Linux distribution on CD-ROM. Red Hat's Initial Public Offering (IPO) in 1999 made it the first publicly traded software company in the State of North Carolina. Red Hat is headquarter on the Centennial Campus of the North Carolina State University, and has regional offices across the US and Canada, as well as international offices in 12 other countries.

Red Hat's product offerings are a combination of software (which is covered by one or more open source licenses) and services (which are not software, and hence not subject to open source licensing) includes the enterprise-class version of Linux as well as content management, web publishing, system management, and application management and monitoring related software. Furthermore, because security is such a critical issue in government, Red Hat has focused on providing information, software tools, training, and consulting to insure that agencies' Linux systems are secure.

Open Source

The range of open-source software now available is as extensive as any single software company's product line. The software ranges from desktop productivity to satellite imaging analysis, from databases to graphic arts, from operating systems to application servers. There, are over 10,000 open source projects and hundreds of thousands of open source developers. The term “open source software” is commonly understood to mean software that is governed by a license that has been approved by the Open Source Initiative (a non-profit organization) and is consistent with the Open Source Definition (available online at http://opensource.org/docs/definition.php).

Open source software differs from proprietary (or single-source or closed-source) software principally in the fact that not only does the software license grant rights to the source code itself, but also the rights to read, modify, and redistribute the software, with or without additional modifications. Some open source licenses permit redistribution without sources whereas others require redistribution with sources, but software redistributed without sources is no longer open source software (and some open source licenses forbid such a change in licensing).

The Open Source Definition does not require that a given distribution be free of charge, but only that whatever the terms of a given software license, that the software license convey rights to the source code at no charge. Because open source software can be distributed without payment of royalty or license fees, there are literally thousands of outlets for open source software as there are many who wish to distribute software for one reason or another. While Red Hat is based on open source code, its distribution of Linux has been thoroughly tested and has organized support available twenty-four hours a day.

The State of North Carolina's STA restricts the use of “open source” products to only those applications that are commercially distributed and supported.

Red Hat provides a commercially distributed and supported “open source” based operating systems that have been tested, certified, packaged, and are currently supported by their professional services division. Detailed accounts of Red Hat’s testing and certification processes are found in the following sections.

Partners

Red Hat has signed global support and reseller agreements with a number of partners including IBM, HP, Dell, Sun Microsystems and Fujitsu. These partnership agreements permit original equipment manufacturers (OEM) to preload Red Hat Enterprise Linux products and sell them as part of a hardware offering. Customers can choose whether to have a direct relationship with Red Hat or an indirect relationship via the OEM partner. Typically customers who want to deploy Linux in only a single context for a single purpose purchase from the OEM; whereas, customers who want to have an enterprise-wide standard that is supported across various hardware platforms (either at the outset or in the future) often choose Red Hat as the primary support provider.

Red Hat has an independent software vendor (ISV) partnership program to certify enterprise offerings from companies such as Oracle, IBM Software, BEA, Veritas, PeopleSoft, Sybase and others. The Red Hat certification program ensures not only that a single enterprise application works with Red Hat Enterprise Linux, but also that it can be expected to work in conjunction with other certified applications. This is important because many enterprise software packages are designed to work with other packages (for example using a BEA web server with an Oracle database and a Veritas volume manager). There are presently more than 50 ISVs who have certified or published plans to certify their software on Red Hat's Enterprise Linux platform.

Information and a listing of all of Red Hat’s Alliance partners – both OEM and ISV – can be found at http://www.redhat.com/partners.

Support Mechanisms

Point of Contact

Red Hat (State of North Carolina) – Red Hat has an Account Manager for the State of North Carolina. Account Managers should be viewed as key strategic resources for Red Hat products and services. They focus on fostering a positive relationship with the State, and can leverage resources as well as partners for guidance and support for Red Hat Linux operating system (OS) based solutions across the enterprise.

As of this publication, the Red Hat Account Executive assigned to the State of North Carolina is Brian Cole (919-754-3700).

Detailed information on Red Hat's solution offerings can be found at http://www.redhat.com/support/sla/na.html.

Telephone numbers for support can be found at http://www.redhat.com/support/techsupport/production/GSS_phone.html.

Books and Magazines

Red Hat works with a variety of publishers to create books and users guides for their products. These publishers include O'Reily as well as Sams and Wiley. Most of the books are based on Red Hat Linux, the non-enterprise version of their product line. The emphasis of this implementation guide, however, is on the enterprise solutions. Nonetheless, these books can be helpful. Good titles include the following:

- ❑ Red Hat Linux Network and System Administration
- ❑ Red Hat Linux Security and Optimization
- ❑ Red Hat RPM Guide
- ❑ Red Hat Certified Engineer Linux Study Guide

A comprehensive list of additional publications can be found through several on-line bookstores.

Red Hat also has a monthly newsletter that interested parties and current users can join to receive additional information about Red Hat products and support. Sign up for newsletters at www.redhat.com/solutions/industries/government.

Learning Services

Having staff who are trained and knowledgeable of Red Hat Enterprise Network, Red Hat Enterprise Linux, and accompanying solutions is an important element of ensuring a smooth, uniform, consistent, and cost effective implementation. Such personnel can be a great support resource for the State. Red Hat has two major certification programs: Red Hat Certified Engineer (RHCE) and Red Hat Certified Technician (RHCT).

The RHCE is aimed at those employees who already possess significant systems administration experience and knowledge in a Unix or Linux environment, and who may want to focus training and certification on Red Hat Linux.

RHCT is an ideal technician level credential for employees supporting Linux systems throughout an agency. RHCT is a good choice for those transitioning to Linux from non-UNIX OSes, or who want to prove their competencies at a midway point on the way to RHCE.

More information on training can be found at http://www.redhat.com/training.

Red Hat Enterprise Network (RHEN)

RHEN is a comprehensive framework allowing you to easily manage your agency's infrastructure - from a few systems to thousands. Through their graphical interface, you can receive automatic software updates, maintain your Red Hat Linux systems, and monitor your heterogeneous environment.

RHEN has three types of architecture:

1. Hosted model - The agency's systems connect with RHEN via the Internet and exchange packages and information from the central RHEN server(s).

2. Hosted/On-site mix model - The agency's systems connect with RHEN for some updates, authentication, or other information. Other aspects of the functionality are delivered directly from application servers running on the customer site.

3. On site model - The entire RHEN operation is run on the agency's local network. In these cases, RHEN also allows for agencies to take their entire networks offline (from the public internet) and still receive full RHEN functionality by running everything on a local network.

Agencies may choose to monitor their networks themselves, or subscribe to Red Hat Enterprise Network Monitoring Module. The monitoring module provides integrated network, systems, application and transaction monitoring, complete with reporting and notification functionality.

A schema of how Red Hat Enterprise Networks operates is as follows:

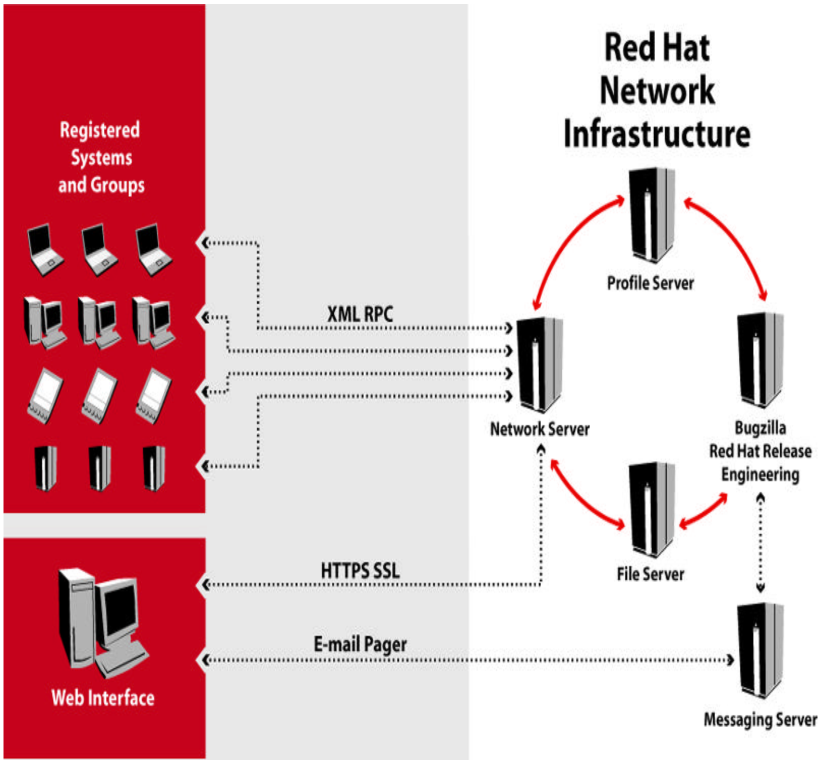

Additional RHEN information can be found at the following site http://www.redhat.com/software/rhen/.

Red Hat Enterprise Linux (RHEL)

The Red Hat Enterprise Linux (RHEL) family is a comprehensive suite of Linux operating systems--designed for mission-critical enterprise computing and certified by top enterprise software vendors

An enterprise class platform must be able to provide not only stability, but also RASM (Reliability, Availability, Scalability, Manageability) Features of the RHEL are listed below.

Reliability: To ensure that RHEL software components are of high quality, Red Hat engineering subjected it to stringent qualification and testing over an extended period of time. Also, extensive stress testing and quality assurance was performed using major enterprise applications and large system configurations. Longer testing schedules and close focus on those operating system features provide for a solid enterprise application.

Availability: RHEL AS provides two integrated high availability technologies: The first, Red Hat Cluster Manager, provides high availability by using a technology widely used by other operating systems – application failover. Red Hat Cluster Manager is discussed further in the next section.

The second high availability technology, IP Load Balancing, provides network load balancing for environments such as web server farms. IP Load Balancing will distribute incoming network requests across a group of servers, which then service the request. Load balancing yields improved performance and, if a single server fails, incoming requests will continue to be redistributed across the remaining servers.

Scalability: The potential performance of larger systems cannot be fully realized with the standard Linux kernel. RHEL products are delivered with kernels that are optimized for SMP systems (up to 8-way). Red Hat has provided several significant kernel enhancements in RHEL that ensure that applications can scale linearly to maximize CPU throughput and processing power.

- *Asynchronous I/O Support* – Allows a process to issue an I/O and immediately continue processing. The application is notified of I/O completion by 1) a process level software interrupt and/or by regular polling of an event flag.

- *I/O Spinlock Contention and Reduction* – Coordinate access to critical kernel code in SMP systems. As a result, different CPUs can simultaneously handle different kernel functions, increasing the kernel performance of the system.

- *Improved Process Scheduler* – RHEL scheduler enhancements schedule a process on the same CPU whenever possible. This means that data and instructions held in the CPU cache will continue to be available for the process each time it executes – greatly improving performance.

- *Bounce Buffer Elimination* – Depending on the specific I/O operation, the enhancement will drastically reduce or eliminate buffer copying while continuing to allow device drivers to access contents of an I/O buffer regardless of its location in memory. RHEL uses a section of kernel virtual address space, whose virtual pages can be used to map any physical address.

Manageability: RHEL products provide enhanced capabilities for monitoring systems and resolving operating system kernel level problems. These features are essential for delivering enterprise-level system support services and are not well developed in standard consumer Linux products.

- *Network Console* – Allows all kernel messages to be logged by another system on the network.

- *Netcrashdump* – Allows a RHEL system to transmit a complete crash dump across the network to a sink node.

Further features of RHEL can be found at the following site www.redhat.com/software/rhel/features.

Products

Red Hat Enterprise Linux ES (RHEL ES)

Red Hat Enterprise Linux ES provides the core operating system and networking infrastructure for a wide range of entry-level and departmental server applications. It is suited for network, file, print, mail, Web, and custom or packaged business applications. Red Hat Enterprise Linux ES is fully compatible across the Red Hat Enterprise Linux family, providing the stability, performance, and support needed for critical application deployments.

RHEL ES provides the same core capabilities as Red Hat Enterprise Linux AS, it is designed for agencies with smaller systems offering up to two CPUs and 4GB of main memory.

Additional ES information may be found at the following site www.redhat.com/software/rhel/es.

Red Hat Enterprise Linux AS (RHEL AS)

Red Hat Enterprise Linux AS (formerly Red Hat Linux Advanced Server) is the core operating system and infrastructure enterprise Linux solution. Supporting the largest commodity-architecture servers--with up to eight CPUs and 16GB of main memory--and available with Red Hat's highest levels of support. RHEL AS is the ultimate Red Hat solution for large agency and data center servers.

RHEL AS is supported by an extensive range of applications from leading ISVs, and is certified on systems provided by Dell, HP, IBM, and Sun.

Also, RHEL AS is certified by DISA (US Defense Information Systems Agency) as COE (Common Operating Environment) Compliant. It is the only Linux distribution to have received this certification. Thus, offering agencies a secure foundation for running mission critical applications.

Additional AS information may be found at the following site www.redhat.com/software/rhel/as/.

Clustering Technology

In addition to enterprise-focused performance features and security, RHEL AS includes two sophisticated highly available clustering technologies:

High Availability Clustering is provided by the Cluster Manager Feature.

High performance and high availability for network infrastructures is provided by the IP Load Balancing (Piranha) feature.

Cluster Manager

The Red Hat Enterprise Linux AS Cluster Manager feature delivers an application failover infrastructure that can be used by a wide range of applications, including:

- ❑ Most custom and mainstream commercial applications
- ❑ File and print serving
- ❑ Databases and database applications
- ❑ Messaging applications
- ❑ Internet and open source application

With Cluster Manager, these applications can be deployed in highly available configurations so that they are always operational--bringing "scale-out" capabilities to enterprise Linux deployments.

For high-volume open-source applications, such as NFS, Samba and Apache, Cluster Manager provides a complete ready-to-use failover solution. For most other applications customers create custom failover scripts using provided templates. Red Hat Professional Services can provide custom Cluster Manager deployment services where required.

Cluster Manager Deployment Tips

- ❑ For Oracle installation, load the oracle software of each cluster server and add keep the database on the shared drive.

- ❑ Use monitoring or a similar service to check the status of the network interface so services can be relocated off a server with a faulty network connection to the users (the cluster does not provide this functionality).
- ❑ For Oracle, subdivide cluster services to disk services and startup script services so as to be able to manage the disk data after shutting down the database.
- ❑ When there is possibility of very high-shared disk access, change the cluster ping interval setting from the default of 2 to 4 seconds.

Cluster Manager Features

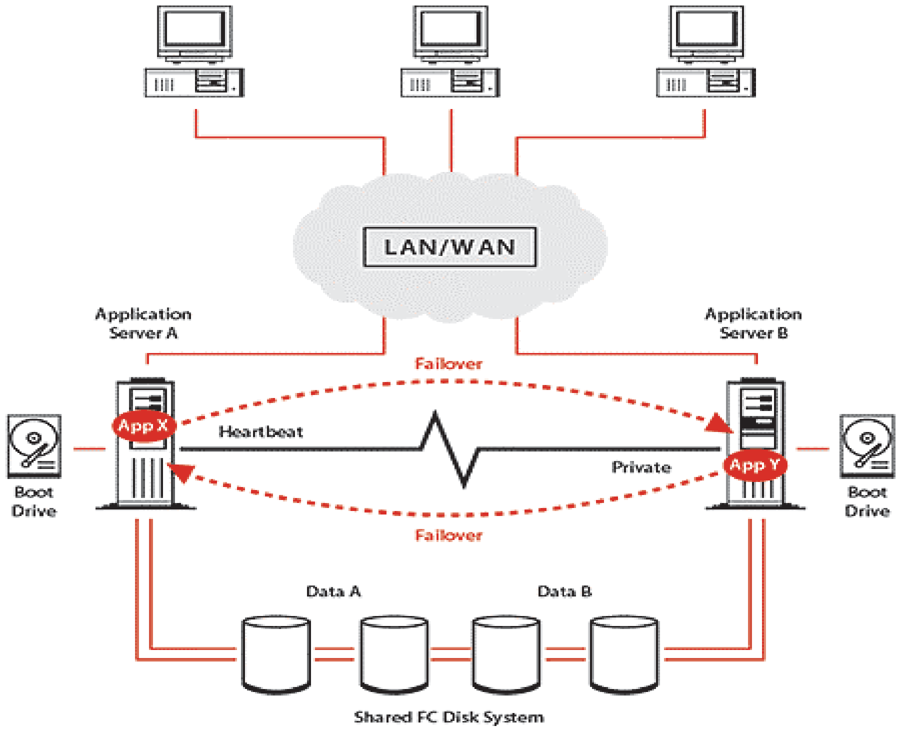

Both servers have redundant connections to disk system(s), but Red Hat Linux Cluster Manager controls access. One server talks to each partition at a time

- ❑ NFS/CIFS Failover: Supports highly available file serving in Unix and Windows environments.
- ❑ Fully Shared Storage Subsystem: All cluster members have access to the same storage.
- ❑ Comprehensive Data Integrity guarantees: Using the latest I/O barrier technology, such as programmable power switches and watchdog timers.
- ❑ SCSI and Fibre Channel support Cluster Manager configurations can be deployed using latest SCSI and Fibre Channel technology. Multi-terabyte configurations can quickly be made highly available.

- ❑ Service failover: Cluster Manager not only ensures hardware shutdowns or failures are detected and recovered from automatically, but also will monitor your applications to ensure they are running correctly, and will restart them automatically if they fail.

IP Load Balancing

IP Load Balancing (often known by its project name, Piranha) provides the ability to load-balance incoming IP network requests across a farm of servers. IP Load Balancing is based on Open Source Linux Virtual Server (LVS) technology, with significant Red Hat enhancements.

The following excerpts from the IP Load Balancing documentation provide an overview of its basic operation:

Load-Balancing Clusters Using IP Load Balancing Servers

An IP Load Balancing (IPLB) cluster appears as one server, but in reality, a user from the Web is accessing a group of servers behind a pair of redundant IPLB routers. An IPLB cluster consists of at least two layers. The first layer is composed of a pair of similarly configured Red Hat Enterprise Linux AS systems. One of these nodes acts as the active IPLB router (the other acts as a backup), directing requests from the Internet to the second layer--a pool of servers called real servers. The real servers provide the critical services to the end-user while the LVS router balances the load to these servers.

The active router serves two roles in the cluster:

1. To balance the load on the real servers.
2. To check the integrity of the services on each of the real servers.

This figure shows a simple IPLB cluster consisting of two layers. On the first layer are two IPLB routers. The active system directs traffic from the Internet to a variable number of real servers on the second layer, which in turn provide the necessary services.

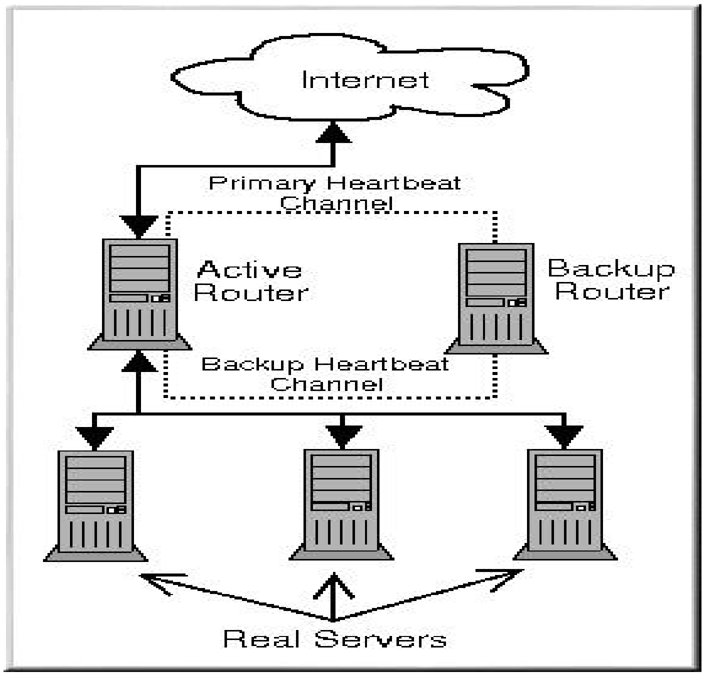

The simple, two-layered configuration used in the figure is best for clusters serving data that does not change very frequently--such as static Web pages--because the individual real servers do not automatically synchronize data between each node. If high-availability shared data access and update is required, a third layer is added to the topology. RHEL AS Cluster Manager is ideal for this purpose, so both clustering types can be used together to create, for example, a high availability 3-tier solution with full transactional capabilities.

IP Load Balancing Deployment Tips

- If persistence settings are used, make sure the quiescence server setting is set to "no."
- Use the firewall mark feature when setting up an HTTP and HTTPS server, to group the two ports together on one virtual server.
- Iptables should be used to perform static NAT translations for DNS access outside the cluster.
- If the web content is non-static, and there are user access time concerns, use the NFS service on a clustered set to provide a central location for the web content. The NFS server should be behind a firewall that allows access to the NFS mounts from the web server and no other hosts. This firewall can be implemented on the NFS server using IPTABLES and can also be a Linux server.

Migration Methodology

Unix and Windows to Linux Standard Methodology

The following methodology should be used to ensure high quality deployment when migrating from Unix or Windows to Linux. If agencies should require assistance, Red Hat's professional services team is available to assist with the migration.

A Unix or Windows to Linux migration should utilize a four-step process to ensure high quality deployment during a migration to Linux.

1. Assessment

The first step in a Linux migration is to identify areas within your infrastructure that can benefit from Linux-based solutions. One's assessment process begins with a thorough review of the existing IT infrastructure, current and future growth estimates, and strategic initiatives.

From there, one needs to develop a Migration Assessment Plan that is used as a foundation for the process. This plan should include the following:

- ❑ A thorough baseline analysis of current infrastructure or application domain, including price/cost, performance benchmarks and functional requirements, compared to the future state based on Linux.
- ❑ An adequate training plan for agency staff.
- ❑ Estimates for total migration costs, risks and a high-level project plan.

2. Planning and Design

The success of a migration to Linux is heavily dependent on thorough planning and design. The Planning and Design phase of a Linux Migration and Integration project evaluates the data gathered in the assessment phase and creates an actionable, comprehensive Implementation Plan that forms the roadmap for the deployment.

The Implementation Plan defines the sequential steps for a successful migration, including the time frame, team structure/resource allocation, key milestones, dependencies, and critical deadlines. A test plan is defined that details testing scenarios, acceptance criteria, and the validation acceptance group made up of client personnel. Additionally, the plan provides a technology baseline for the Linux infrastructure deployment and takes into account services and applications that need to be deployed on the Linux platform, along with recommended hardware configurations to meet performance requirements.

The Implementation Plan also documents recovery and contingency plans to mitigate risk during transition to the new production environment.

3. Development and Validation

The Development and Validation phase of a Linux Migration and Integration includes any required development, including core build development, system deployment architecture, porting, performance tuning and code optimization, as well as security hardening, testing and validation of the future environment.

This phase requires a coordinated effort to ensure that the deployment does not disrupt on-going business operations. The result of the Development and Validation phase is a fully tested and functional solution that is ready for deployment.

4. Deployment and Operations

Open source technology deployments should be supported by proven business processes and skilled administrators, in order to effectively manage the solution once it is moved into production. The constant challenge to innovate and improve operational processes is often impeded by a lack of skilled technical resources in this area. This need to constantly upgrade and deploy cutting edge technologies puts a tremendous amount of strain on the resources within the IT department.

Red Hat's Linux Deployment program delivers core support and maintenance services that will help you streamline and manage your open source infrastructure in a more efficient manner. Red Hat is qualified to provide knowledge transfer to IT organizations, helping them to continually improve their production environments. Red Hat's process driven approach to professional service delivery can also help you assess and improve the internal processes that govern open source infrastructure deployment and management.

Red Hat Enterprise Linux OS Installation

North Carolina Government relies upon a complex communication infrastructure and connectivity via public and private networks. This places an ever-increasing demand on agencies to protect and control access to information across their networks. The following guidelines will not only provide a more secure operating environment, it will allow agencies to improve their operational efficiency and effectiveness.

Installation Planning Process

A full installation of Red Hat Linux contains up to 1200 application and library packages. It is important to note that a default installation of the Red Hat Enterprise Linux is not considered to be hardened. System administrators should not opt to install every single package in the distribution, preferring instead to follow the industry best practice of installing a base set of packages that include only the required server applications.

A common occurrence among system administrators is to install the operating system without paying attention to what programs are actually being installed. This can be problematic because unneeded services might be installed, configured with the default settings, and possibly turned on by default. This can cause unwanted services, such as Telnet, DHCP, or DNS to be running on a server without the administrator realizing it, which in turn can cause unwanted traffic to the server, or even, a potential pathway for crackers.

With a detailed OS Installation Plan a system administrator is more likely to successfully install the OS with consistency and accuracy. Furthermore, consistency is a key factor in security. A consistent installation plan will make it easier for you to maintain secure configurations and help you to identify security problems.

Documentation

When beginning the OS Installation Plan, begin by documenting how the server will be used. Consider the following:

- ❑ What categories of information will be stored on the server?
- ❑ What categories of information will be processed on the server?
- ❑ What are the security requirements for that information?
- ❑ What network services will be provided by the server?
- ❑ What are the security requirements for those services?

After documenting the answers to the above questions, and determined the minimal set of services and applications, then the system administrator needs to ensure that only those are installed on the host. Following the steps listed hereafter will help ensure consistency and accuracy.

Identify the users or categories of users of the computer.

Document the categories of users that will be allowed access to the provided services. For public servers connected to the Internet, the category of users is probably everyone. For internal servers, categorize users by their organizational department, physical location, or job responsibilities. Additionally, develop a category of administrative users who will need access to administer the network server and possibly another category for backup operators.

Restrict access to servers to only those system administrators responsible for operating and maintaining the server.

In the procedures, include steps to implement all the decisions made in the steps above and describe all the parameters that are set during installation.

In many cases, the parameters are recorded in scripts or configuration files that are executed or read during various phases of the installation. Make all parameter choices explicit, even if they match Red Hat's current default settings. (This may seem to be unnecessary, but it can prevent security problems if subsequent reuse of scripts or configuration files to configure additional servers are used.) Explicit choices will still be used even if Red Hat's defaults change with new releases. The installation procedure should also specify Red Hat's security-related updates or patches that are to be applied to the operating system.

If possible, have a single person perform the installation procedure for each server and capture each installation step in a documented manner (such as through using a checklist). Then, have another system administrator perform the installation on other servers using the documented procedures, thus, validating the procedures.

Installation Location

Systems are often times susceptible to attack during the initial implementation of the OS. Until the OS has been properly configured and patched, existing security vulnerabilities maybe exploited. Adhering to the following guidelines will help ensure a successful OS installation.

Configure the host system BIOS to boot only from a hard drive. Set a BIOS administration password to prevent changing the boot sequence. Also set a boot password unless the serve must be rebooted remotely.

Install the OS on the host in a physically secure environment to minimize exposure to attack and other undesired activity. Also, no one other than the system administrator responsible for installing the OS should have physical access to the system while it is being installed and made ready for deployment. Enterprise servers should only be deployed in controlled environments such as a data center. As a general rule, do not deploy servers in an individual's office. Additionally, locate the servers so unauthorized viewing of the monitor and keyboard cannot occur.

If you must install the OS on a host connected to a network, then that network should be separate from all production networks and disconnected from public networks, including the Internet until the installation is complete. Also, the network should not have any additional systems or devices attached to it that are not explicitly required for installations. Moreover, no one other than authorized personnel responsible for the installation of the host should have access to the isolated network, system, or devices attached to it.

For security purposes, ensure that the network cabling is not placed in a physical location where it can be easily accessed. Note that this requires a trade off between the convenience of network access for network maintenance and security.

OS Configuration

Red Hat Linux provides the capability to specify access privileges individually for directories, files, devices, and other objects. By carefully setting and documenting access controls, the system administrator can reduce both intentional and unintentional security breaches. To help ensure a successful OS configuration, the system administrator should adhere to the following guidelines.

- ❑ Only install the services required for the host to perform its required functions.
- ❑ Configure GRUB or LILO with a password to prevent users from changing boot options, such as runlevel.
- ❑ Perform an initial Port Scan using a sniffer tool to ensure that only the desired ports are open for the IP-based services intended to be exposed.
- ❑ Utilize xinetd, PAM and IPTABLES as appropriate to limit external access to networked machines to only those services required on the host.
- ❑ After installation, check the running services and stop any that are running that are not required, and use chkconfig to ensure that they do not restart on reboot.

Deny read access to protect confidentiality of information, and deny unnecessary write access, this can help maintain the integrity of information.

Limit the execution privilege of system-related tools to authorized system administrators. This can prevent most users from making configuration changes that could reduce security. It also can restrict the ability of intruders to use those tools to attack the system or other systems on the network.

As previously stated, implement and document access controls during initial installation and configuration of the operating system, and carefully monitor and maintain them afterwards.

Identify the protection needed for directories, files, devices, and objects. This is best done by constructing a matrix with categories of files and objects on one axis and groups of users defined by roles and access authority on the other. Then record in the matrix the kinds of access privileges allowed for that class of objects and that class of users. The privileges are based on the security requirements, such as confidentiality, integrity, and availability of the various classes of resources. One might also want to distinguish local access privileges from network access privileges for a class of files.

When you take the previous step, one may identify categories of users not sufficiently detailed in the computer deployment plan. Configure the operating system to recognize the needed user groups, and then assign individual users to the appropriate groups.

Additionally, configure access controls for all protected directories, files, devices, and other objects, using the previously created matrix as a guide. Every change or decision not to change each object's permission should be documented along with the rationale.

Disable write/modify access permissions for all executable and binary files.

Restrict access of operating system source files, configuration files, and their directories to authorized administrators.

Limit world-writable files unless required by application program.

Mount file systems as read only and nosuid to preclude unauthorized changes to files and programs.

Assign an access permission of immutable to all kernel files

Establish all log files as "append only" if that option is available.

Review logs frequently to identify suspicious entries. Logwatch or other similar log analysis tools can assist in identification of these entries.

Configure the OS so that newly created files and directories inherit appropriate access controls, and that access controls propagate down the directory hierarchies as intended when you assign them.

Use SSH for remote access instead of telnet, ftp or any of the r-commands. SSH provides a secure and encrypted communications path between hosts, as well as SCP for secure file transfers.

For additional information regarding the Platform Domain, visit the Statewide Technical Architecture website at http://ets.state.nc.us/NCSTA/ets_index.html

Credits

Participating Agencies

1. **Office of Information Technology Services**

2. **Office of Enterprise Technology Strategies**

3. **Department of Health and Human Services**

4. **Department of Environment and Natural Resources**